

From “Is it unconfounded?” to “How much confounding would it take?”: Applying the sensitivity-based approach to assess causes of support for peace in Colombia

Short Title: How much confounding would it take?

Abstract

Attention to the credibility of causal claims has increased tremendously in recent years. When relying on observational data, debate often centers on whether investigators have ruled out *any* bias due to confounding. However, the relevant scientific question is generally not whether bias is precisely zero, but whether it is problematic enough to alter one’s research conclusion. We argue that sensitivity analyses would improve research practice by showing how results would change under plausible degrees of confounding, or equivalently, by revealing what one must argue about the strength of confounding to sustain a research conclusion. This would improve scrutiny of studies in which non-zero bias is expected, and of those where authors argue for zero bias but results may be fragile to confounding too weak to be ruled out. We illustrate this using off-the-shelf sensitivity tools to examine two potential influences on support for the FARC peace agreement in Colombia.

Keywords: sensitivity analysis, observational studies, causal inference, partial identification, FARC

Chad Hazlett

Associate Professor, Political Science and Statistics
University of California-Los Angeles
chazlett@ucla.edu

Francesca Parente

Assistant Professor, Political Science
Christopher Newport University
francesca.parente@cnu.edu

Supplementary material for this article is available in the appendix in the online edition. Replication files are available in the JOP Dataverse (<https://dataverse.harvard.edu/dataverse/jop>). The empirical analysis has been successfully replicated by the JOP replication analyst.

1 Introduction

Many important research questions cannot be answered by randomized experiments, whether because researchers cannot ethically or practically randomize the treatments involved, or because they would like to study the effect of real world events that have already occurred. Observational alternatives to randomized experiments involve demanding and typically untestable assumptions. The most commonly employed strategies rely on an assumption of “no unobserved confounding” and then adjust for observables using estimators such as stratification, regression, matching, or various weighting approaches. Other strategies that seek to limit confounding biases, such as difference-in-difference, regression discontinuity, instrumental variables, and quasi- or natural-experiments, offer opportunities to improve inference, but are not always available and require their own demanding, largely untestable assumptions.

Fortunately, recent increases in attention to the credibility of observational findings have heightened many researchers’ awareness of the risks posed by confounding. In studies where investigators adjust for observables, debate typically centers on whether *any* unobserved confounding may remain. Such “all or nothing” arguments about the threat of confounding have been a useful antidote to practices of turning a blind eye to confounding. However, we argue that this approach is suboptimal and sometimes problematic as a means of judging the risk that a result is driven by confounding. First, when authors argue for a total absence of unobserved confounding, readers and critics can and should be able to disagree, e.g. by proposing unobserved confounders that have not been ruled out. A means of debating what damage would be done by such confounders is thus required. Second and more fundamentally, the scientifically relevant concern is not usually “is there exactly zero confounding bias?”, but rather “might confounding have altered the substantive research conclusion?” Yet, the arguments deployed to claim that an observational study is confounding-free are typically silent on the question of how fragile the result is in the face of possible errors in that judgment.

Consequently, even a study that is arguably confounding-free (e.g. on the basis of a convincing natural experiment) could be sensitive to very weak confounders—perhaps too small to be definitively ruled out once confronted with this information. On the other hand, studies in which the authors recognize confounding may exist could in fact require powerful confounding to alter their conclusions—perhaps so powerful that such confounding can be deemed implausible. Many studies will fall in between these extremes. In all cases, reporting how much confounding

it would take to overturn a given result is transparency-enhancing. While conventional statistics like the t-statistic, coefficient, or p-value cannot correctly answer this question, easy-to-compute and interpret “sensitivity statistics” can (Cinelli and Hazlett, 2020).

We employ off-the-shelf sensitivity analysis tools to demonstrate the transition from “all or nothing” claims about confounding to practices that reveal and communicate how much confounding it would take to change the research conclusion. For concreteness, we focus on an application examining two proposed influences on civilian voting behavior in the 2016 peace deal with the FARC in Colombia. This question has attracted great scholarly interest, with at least 14 observational papers since 2017 (see Appendix A for an expanded review). Scholars seeking to explain this real-world event have coalesced around two potential influences on voting behavior: (1) *exposure to violence* and (2) *political affiliation* with the deal’s champion, President Santos. Their findings have been remarkably consistent in supporting both the exposure to violence (e.g. Tellez, 2019) and political affiliation hypotheses (e.g. Krause, 2017). Yet, none of these papers has been positioned to argue that their result is due to a causal effect rather than confounding.¹ We find that even very weak confounders could explain away the relationship between exposure to violence and FARC support. However, the effect of political affiliation requires very large confounding to overturn. Such analyses do not settle the issue of causality, but instead transparently communicate what assumptions about the strength of confounding would need to hold to achieve a given result, enabling more reasoned and productive debate than the “all or nothing” approach alone.

2 The sensitivity-based approach

We use the term sensitivity analysis specifically to refer to statistical analyses that characterize how a result changes under postulated unobserved confounding. Such tools have a long history dating back to at least Cornfield et al. (1959). Yet they remain largely unutilized: In the top three general interest political science journals in 2019, only 7 out of 119 (5.9%) papers whose conclusions may be altered by unobserved confounding included formal analyses describing the strength of confounding required to change the conclusion (see Appendix B).

Fortunately, sensitivity statistics can be computed that correctly characterize how an estimate

¹Experimental alternatives that seek to test particular claims are of course possible; see, e.g. Matanock and Garbiras-Díaz, 2018. However, our interest here lies in practices of making causal claims from observed, non-experimental data.

would change under varying degrees of postulated confounding, or equivalently, the strength of confounding required to achieve a given change in an estimate. Any sensitivity analysis approach can be useful if it (i) correctly and clearly illuminates the strength of confounding required to alter the research conclusion, and (ii) aids researchers, as cogently as possible, in understanding that strength of confounding or in debating whether such confounding plausibly exists. While similar analyses could and should be conducted for other estimation approaches (e.g. matching and weighting estimators), we focus here on the sensitivity of linear regression models because (i) linear regression is ubiquitous in political science (including in the works we consider in this application); and (ii) because the linearity of these approaches proves especially convenient for analyzing sensitivity. Of the sensitivity toolkits developed for linear regression, we employ [Cinelli and Hazlett \(2020\)](#), which elaborates on the concept of omitted variables bias that will be familiar to many readers. This approach also conveniently produces summary statistics to characterize sensitivity, and offers rigorous tools for comparing hypothetical confounders to observed benchmarks. We discuss the broader landscape of sensitivity approaches and how they compare in Appendix B.

While space prevents a full review of sensitivity analysis, we provide here the minimal technical material required to discuss and responsibly employ these tools. We refer readers interested in the technical details to [Cinelli and Hazlett \(2020\)](#) as well as to Appendix C, where we discuss finer points and commonly raised questions and concerns regarding these tools.

The key to understanding this approach begins with considering “the regression you wish you ran”. Suppose that the regression of an outcome (Y) on the treatment (D) would be biased by many possible confounders, but those confounders would be controlled for if we included covariates X and Z (both containing one or more variables). In other words, controlling for X and Z together would render the D - Y relationship unconfounded. It would be desirable to then run the regression

$$Y = \hat{\tau}D + \mathbf{X}\hat{\beta} + \hat{\gamma}Z + \hat{\epsilon}_{\text{full}} \tag{1}$$

to approximate the conditioning on X and Z and provide a linear summary of the remaining relationship between D and Y . Unfortunately, while the variables in X are observed, those in Z are unobserved. The regression actually run is instead the “restricted” regression,

$$Y = \hat{\tau}_{\text{res}}D + \mathbf{X}\hat{\beta}_{\text{res}} + \hat{\epsilon}_{\text{res}}. \tag{2}$$

The central question is how the observed estimate ($\hat{\tau}_{\text{res}}$) differs from the desired one, $\hat{\tau}$. We thus define $\widehat{\text{bias}} := \hat{\tau}_{\text{res}} - \hat{\tau}$, the difference between the estimate *actually* obtained and what would have been obtained in the same sample had the missing covariate(s) Z been included.

As shown in Cinelli and Hazlett (2020), the bias due to omission of Z can be written as

$$|\widehat{\text{bias}}| = \text{se}(\hat{\tau}_{\text{res}}) \sqrt{\frac{R_{Y \sim Z | \mathbf{X}, D}^2 R_{D \sim Z | \mathbf{X}}^2}{1 - R_{D \sim Z | \mathbf{X}}^2}} (\text{df}), \quad (3)$$

The two sensitivity parameters here are partial R^2 values, $R_{Y \sim Z | \mathbf{X}, D}^2$ and $R_{D \sim Z | \mathbf{X}}^2$. The first ($R_{Y \sim Z | \mathbf{X}, D}^2$) describes how the fraction of Y *not already (linearly) explained by X and D* can be (linearly) explained by Z . The second, $R_{D \sim Z | \mathbf{X}}^2$ is similarly the fraction of the treatment status explained by confounding, after accounting for the observed covariates. A similar expression is available describing the adjusted standard error (see Cinelli and Hazlett, 2020) in terms of the same parameters, $R_{Y \sim Z | \mathbf{X}, D}^2$ and $R_{D \sim Z | \mathbf{X}}^2$. A fundamental fact conveyed by these formulas is that these two parameters jointly characterize the only properties we need to know about an omitted confounder in order to determine how the point estimate, standard error, t-statistic, or p-value would be changed by including that variable.

To communicate the fragility of a result in the face of unobserved confounding we employ summary statistics described in Cinelli and Hazlett (2020), as detailed in Table 1. Additionally, “benchmarking” tools provide one useful way to interpret sensitivity analyses and argue for bound on confounding, by comparing the strength of confounding required to change the result to the explanatory power of one or more observed covariates. This aids, first, in understanding the magnitude of confounding required to change an answer by restating it in terms of observed covariates, for which we have stronger intuitions regarding the strength of relationship with treatment and outcome. Further, in some cases users may employ their domain knowledge and information about treatment assignment to argue that unobserved confounding is not likely to explain “ k times more of the treatment assignment and outcome” than a given observable. If confounding of such strength would not change the conclusion, this can be compelling evidence for the credibility of the research conclusion. Or, if one finds that even an optimistic assumption would still allow enough confounding to change the research conclusions, we know the original conclusion would be difficult to satisfactorily defend.

We employ these tools using the `sensemkr` package for R. Replication code for our analyses can

Table 1: Sensitivity summary statistics from [Cinelli and Hazlett \(2020\)](#)

Statistic	Interpretation
$R_{Y \sim D \mathbf{X}}^2$	The partial R^2 of the treatment with the outcome: In the extreme case where confounders explain 100% of the residual variance of the outcome, this tells us how much of the residual variance in the treatment such confounders would need to explain to bring the estimated effect down to zero.
RV	The robustness value: Confounding that explains at least $RV\%$ of residual variance in the treatment and in the outcome would reduce the implied estimate to zero. If both $R_{Y \sim Z \mathbf{X}, D}^2$ and $R_{D \sim Z \mathbf{X}}^2$ are less than the RV , confounding is not sufficient to eliminate the effect.
$RV_{q, \alpha}$	The amount of confounding required such that an effect estimate reduced by the fraction q (e.g. 50%) would fall just within the confidence interval defined by α , i.e. the 95% confidence interval for $\alpha = 0.05$.

be found at <https://dataverse.harvard.edu/dataverse/jop>. Additional examples, tutorials, and resources can be found in Appendix C.

3 Support for the FARC peace agreement

We now illustrate these tools by applying them to examine different explanations for public support of the 2016 peace deal with the FARC in Colombia. In October 2016, Colombians voted in a referendum on a peace agreement with the FARC, a leftist guerilla group. The peace deal was ultimately rejected, but given the immense variation in municipality-level vote share in favor of the deal (ranging from 19% to 96% in towns with at least 1,000 voters), many have sought to explain levels of support for the deal. Scholarship has coalesced around two explanations: exposure to FARC violence and political affiliation with President Santos (see Appendix A). As is often the case for substantively important questions in political science, the “treatment” of interest (exposure to violence or political affiliation) cannot be assigned randomly, nor can we argue that it is as-if random conditional on some set of observables X . Nevertheless, previous papers on this topic have relied upon various covariate adjustment approaches to address confounding (e.g. regression, matching, and weighting). Unfortunately, none of these can hope to argue for an absence of unobserved confounders, so the results must be understood as potentially biased.²

²Concretely, one troubling example of a potential confounder we cannot observe is “latent sympathy for the FARC”. For example, suppose that those more sympathetic to the FARC also tend to be more supportive of one party or leader and more supportive of the deal, while perhaps living in areas that the FARC refrain from attacking. Such “common cause confounding” would confound both the violence and political affiliation accounts in the observed directions.

3.1 Assessing evidence for the effect of exposure to violence

We consider first a naive, direct comparison by regressing $Deaths_{i,2011-2015}$, the number of deaths in municipality i committed by the FARC between 2011 and 2015, on the proportion voting “Yes” in municipality i . Table 2 presents results for such a regression (Model 1), together with the sensitivity quantities. The coefficient of 1.45 ($p < 0.001$) on violence in 2011-2015 implies that with each additional observed death, we expect to see a 1.45 percentage point higher level support for the FARC peace deal. In usual research practice, despite an awareness that confounders may have generated an unknown amount and direction of bias, such a result is typically communicated as “suggestive evidence” that exposure to violence causes higher support for peace.

Table 2: Augmented regression results for violence

Outcome: <i>Vote for peace deal</i>							
Treatment:	Est.	SE	t-stat	$R^2_{Y \sim D \mathbf{X}}$	RV	$RV_{\alpha=0.05}$	df
1. <i>Deaths 2011-2015</i>	1.45	0.30	4.90	2.1%	14%	8.4%	1121
2. <i>Deaths 2011-2015</i>	0.61	0.29	2.11	0.4%	6.1%	0.4%	1115


The sensitivity quantities added to Table 2 quickly characterize how strong confounding would have to be to alter our conclusion. The robustness value (RV) tells us that confounding that explains at least 14% of the residual variance in both violence and support for peace would be enough to eliminate this effect entirely. Recalling that taking the square root of an R^2 allows interpretation on the usual correlation scale, this means a hypothetical confounder with a (residual) correlation of 0.34 to both the treatment (violence) and outcome (support for the peace deal) would be sufficient to explain away the entire result. Similarly, the $RV_{\alpha=0.05}$ value tells us that confounding explaining 8.4% of residual variance in violence and support would reduce the estimate to the edge of statistical insignificance. Finally, the $R^2_{Y \sim D | \mathbf{X}}$ tells us that if an unobserved confounder explains 100% of the remaining outcome variation, such a confounder would have to explain only 2.1% of the residual variation in the violence treatment in order to reduce the estimated effect to zero.

In practice, researchers typically include control variables in hopes of mitigating uncontrolled confounding. Model 2 includes a number of potentially worrying observed confounders (see Appendix D for details). The resulting coefficient estimate is 0.61 ($t=2.11$). While this may pass the bar for publication as suggestive (and statistically significant) evidence in some settings, we emphasize that the statistical significance of a result alone says nothing about vulnerability in

the face of unobserved confounding. A confounder explaining only 6.1% of the residual variation in both violence and support for peace would eliminate the effect; a confounder explaining only 0.4% of both would reduce the effect to the boundary of insignificance at the $\alpha = 0.05$ level. We expect readers would not have trouble coming up with confounders that arguably relate to treatment and outcome at least as strongly, such as sympathy for the FARC.

Beyond these summary statistics, a more complete picture entails visualizing the bias as we vary the strength of postulated confounding (Figure 1). These plots also visualize the results of benchmarking examples, showing how bias would be bounded subject to assumptions that relate the strength of unobserved confounding to the explanatory power of one or more observables. For example, let us assume that *political affiliation* is “stronger” than confounding, in the sense that it explains a greater share of both the outcome and treatment than can true confounding. We proxy for political affiliation using vote share for Santos in 2010, to ensure it is pre-treatment with respect to 2011–2015 violence. While assuming such a limitation on confounding may already be indefensible—it is hard to argue that *political affiliation* explains more remaining variation in exposure to violence than all other confounders could—it already admits more than enough confounding to change the sign of the estimate. Specifically, the point in Figure 1 marked “1x santos10” indicates what the adjusted estimate would be had confounding of this strength been present, bringing the original estimate (+0.61) to the opposite side of zero (-0.24).

Figure 1: Contour plot showing sensitivity to hypothesized confounding.



Finally, sensitivity analyses can aid readers and reviewers even when authors did not originally provide these analyses. Using reported results in [Tellez \(2019\)](#) and [Pechenkina and Gamboa \(2022\)](#), we find an RV of 4.7% and 9.7% and an $RV_{\alpha=0.05}$ of just 1.8% and 4%, respectively,

reflecting sensitivity values similar to our own (see Appendix E for details).

3.2 Assessing evidence for the effect of political affiliation

We estimate coefficients in a model regressing the proportion voting “Yes” in municipality i on the municipal vote share for President Santos (2014), including controls for potential confounders (see Appendix D for details). Table 3 shows augmented regression results. The estimated effect of Santos vote share on support for peace (0.67) is positive and statistically significant. Vote share for Santos in 2014 explains 59% ($R^2_{Y \sim D|\mathbf{X}}$) of the residual variation in support for peace, meaning that even confounding that explains 100% of the residual variation in the outcome would need to explain 59% of the residual variation in vote share for Santos in order to eliminate the estimated effect. Confounding that explains less than 68% (RV) of both vote share for Santos and support for peace would not be sufficient to eliminate the effect.³ For the 95% confidence interval to just include zero, confounding would have to explain 66% of residual variance in both treatment and outcome ($RV_{\alpha=0.05}$).⁴ In short, very strong confounding would be required to overturn these results, as we discuss in context below.

Table 3: Augmented regression results for political affiliation

Outcome: <i>Vote for peace deal</i>							
Treatment:	Est.	SE	t-stat	$R^2_{Y \sim D \mathbf{X}}$	RV	$RV_{\alpha=0.05}$	df
<i>3. Santos 2014 vote share</i>	0.67	0.02	37.5	59%	68%	66%	983

Again, where other studies employed OLS we can determine how sensitive their results would be as well. In Krause (2017), the coefficient for Santos 2014 vote share in a similar model is 0.62, close to our estimate. The t-statistic of 45 together with 1,069 residual degrees of freedom would produce an RV of 72%, also close to our own estimate of 68% (Appendix E). We discuss further analyses, including benchmarking, on the political affiliation hypothesis in Appendix F.

³Recall that an R^2 is just the square of r , the correlation coefficient. Thus, a partial R^2 of 0.68 corresponds to a correlation of $\sqrt{68\%} \approx 82\%$ after accounting for the other covariates—an extremely high correlation by any standard.

⁴While it appears that a similar if coarser conclusion could be drawn by simply comparing the t-statistics of the models, this is only because the sample sizes are similar across models here. More generally t-statistics and p-values cannot reflect how strong confounding must be to alter our conclusions without adjustment based on the degrees of freedom. For example, a coefficient with a t-statistic of 10 and only 200 degrees of freedom has an RV of 50%; however with one million degrees of freedom, the same t-statistics correspond to an RV below 1.0%.

4 Discussion

Whereas asking whether a study is *entirely* free of confounding can be useful in drawing attention to the threat posed by confounding, the central concern in observational research is not whether there is *any* confounding, but whether the research conclusion might have been substantively affected by it. In any setting where arguments for exactly zero confounding may fail, we argue that more can be learned and transparently conveyed by reporting how much confounding it would take to substantively alter a research conclusion, and by using tools such as benchmarking to aid in debating whether such confounding can or cannot be readily ruled out.

The example and tools shown here illustrate how a sensitivity-based framework, if widely adopted, could improve how we conduct, communicate, and evaluate observational research. First, this approach suggests standards for empirical research seeking to make causal claims. Summary sensitivity quantities reported in augmented regression tables provide readily interpretable information about sensitivity to unobserved confounding. Here, sensitivity analyses have helped to determine that conclusions regarding the role of exposure to violence are currently too fragile to hold in high confidence. Even under the most favorable model used, confounding that explains just 6.1% of the residual variance of exposure to violence and support for peace would eliminate the result entirely, and a confounder explaining just 0.4% would reduce it below conventional levels of statistical significance. For the political affiliation hypothesis, however, a confounder explaining 100% of the residual variation in support for peace would have to explain 59% of the residual variation in political affiliation to alter our conclusions.

Second, these tools can improve how we judge the credibility and value of research projects by (i) providing readers a way of assessing how susceptible results are to confounding and (ii) encouraging critics to elaborate their concerns about confounding by arguing for specific confounders that they believe are strong enough to make a difference. We must remember that the large robustness value for political affiliation does not rule out the possibility that confounding has altered our conclusion. Yet, given what we know about the degree of confounding required to alter our result, a colleague or reviewer cannot suggest “just any confounder” would be sufficient to meaningfully change our conclusion. Rather, those suggesting a particular confounder are obligated to argue that it could plausibly explain the amount of variation in treatment and outcome required to alter the results.

Ultimately, we hope these tools help to bring about a change in how we value empirical projects

when arguments about the absence of confounding are fallible: A paper is not to be judged by whether it convinced us that the design leaves zero confounding, but rather by how it informs our understanding of results under degrees of confounding that may plausibly exist.

Acknowledgments

We thank Neal Beck, James Bisbee, Matt Blackwell, Youssef Chouhoud, Darin Christensen, Carlos Cinelli, Jeff Gill, Jens Hainmueller, Erin Hartman, Leslie Johns, Andrew Kirkpatrick, Jeff Lewis, Aila Matanock, Matto Mildenerger, Lauren Peritz, Margaret Peters, Michael Ross, Leah Stokes, Juan Tellez, Teppei Yamamoto, and three anonymous reviewers for valuable feedback. Previous versions of this work were presented at meetings of the American Political Science Association, Empirical Studies of Conflict, Society for Political Methodology, and Visions in Methodology and at seminars at UC Santa Barbara and MIT.

References

- Cinelli, Carlos and Chad Hazlett. 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society, Series B* 82(1):39–67.
- Cornfield, Jerome, William Haenszel, E. Cuyler Hammond, Abraham M. Lilienfeld, Michael B. Shimkin and Ernst L. Wynder. 1959. "Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions." *Journal of National Cancer Institute* (23):173–203.
- Krause, Dino. 2017. "Who Wants Peace? The Role of Exposure to Violence in Explaining Public Support for Negotiated Agreements: A Quantitative Analysis of the Colombian Peace Agreement Referendum in 2016." Unpublished Master's Thesis, Uppsala University.
- Matanock, Aila and Natalia Garbiras-Díaz. 2018. "Considering Concessions: A Survey Experiment on the Colombian Peace Process." *Conflict Management & Peace Science* 35(6):637–655.
- Pechenkina, Anna O. and Laura Gamboa. 2022. "Who Undermines the Peace at the Ballot Box? The Case of Colombia." *Terrorism & Political Violence* 34(2):197–217.
- Tellez, Juan Fernando. 2019. "Worlds Apart: Conflict Exposure and Preferences for Peace." *Journal of Conflict Resolution* 63(4):1053–1076.

Biographical Statements

CHAD HAZLETT is an ASSOCIATE PROFESSOR at UNIVERSITY OF CALIFORNIA-LOS ANGELES, LOS ANGELES, CA 90025. FRANCESCA PARENTE is an ASSISTANT PROFESSOR at CHRISTOPHER NEWPORT UNIVERSITY, NEWPORT NEWS, VA 23606.

Supplementary Material for "From 'Is it unconfounded?' to 'How much confounding would it take?': Applying the sensitivity-based approach to assess causes of support for peace in Colombia"¹

Chad Hazlett; Francesca Parente

Appendix A The FARC referendum

In October 2016, Colombians voted in a referendum on a peace agreement with the FARC, a leftist guerilla group. The peace deal was ultimately rejected, but given the immense variation in municipality-level vote share in favor of the deal (ranging from 19% to 96% in towns with at least 1,000 voters), many have sought to explain levels of support for the deal. The surprising defeat of the peace deal generated a wealth of scholarship on the topic (e.g. Krause, 2017; Matanock and Garcia-Sánchez, 2017; Dávalos et al., 2018; Liendo and Braithwaite, 2018; Matanock and Garbiras-Díaz, 2018; Matanock, Garbiras-Díaz and García-Sánchez, 2020; Branton et al., 2019; Gallego et al., 2019; Hurtado-Parrado et al., 2019; Masullo and Morisi, 2019; Meernik, 2019; Pechenkina and Gamboa, 2022; Tellez, 2019a,b; Revelo and Sottilotta, 2020). In the numerous papers published on what influenced voting in the 2016 FARC referendum, scholars have emphasized two possible causes of support: exposure to FARC-related violence and political affiliation with President Santos.

First, scholars have argued that exposure to FARC violence increased support for peace. For example, Tellez (2019b) finds that citizens in municipalities labelled "conflict zones" by the government were more likely to report that they supported the peace process and concessions to FARC in AmericasBarometer surveys. Using a measure of violence per capita at the municipal level, Branton et al. (2019) find that municipalities exposed to more violence were more in favor of the peace deal. Dávalos et al. (2018) find a similar result using cumulative counts of victims of FARC violence and internally displaced persons, as do Pechenkina and Gamboa (2022) using a binary indicator for FARC attacks between 2011 and 2013. These findings are consistent with a vast theoretical and empirical literature on the effects of violence on attitudes and political behavior (Weintraub, Vargas and Flores, 2015; Bauer et al., 2016; Hazlett, 2020; Fabbe, Hazlett and Simmazdemir, 2019, among others). In the Colombian case in particular, Meernik (2019)

¹Replication files are available in the JOP Dataverse <https://dataverse.harvard.edu/dataverse/jop>. The empirical analysis has been successfully replicated by the JOP replication analyst.

and [Hurtado-Parrado et al. \(2019\)](#) emphasize the role that attitudes and emotions play in Colombians' desires for reconciliation and peace.

We note – as did some journalists – that the relationship could also have gone in the opposite direction. For example, according to the BBC, “In ... Casanare... 71.1% voted against the deal. It is an area where farmers and landowners have for years been extorted by the FARC and other illegal groups” ([BBC News, 2016](#)). Again, a theoretical literature would be consistent with this result. For example, [Petersen and Daly \(2010\)](#) stress the role of anger and emotions in determining attitudes toward peace, with exposure to violence making victims less likely to support reconciliation. However, the published results on this topic have almost exclusively focused on the positive relationship between exposure to violence and support for peace.²

Second, scholars have found that greater support for the peace deal is associated with greater support for the deal's champion, President Santos. [Krause \(2017\)](#), [Branton et al. \(2019\)](#) and [Dávalos et al. \(2018\)](#) all find that the municipal vote share for President Santos in the 2014 presidential election is a strong predictor of how the municipality voted on the peace deal in 2016.³ [Masullo and Morisi \(2019\)](#) and [Gallego et al. \(2019\)](#) demonstrate the importance of information and endorsements in influencing opinions on the peace deal. [Revelo and Sottilotta \(2020\)](#) also emphasize the politicization of the peace process and the effect this had on attitudes toward the deal. These findings are all consistent with a broad literature on the power of elites in directing voting behavior or influencing opinions on policies (e.g. [Zaller, 1992](#); [Lupia, 1994](#); [Levendusky, 2010](#); [Nicholson, 2012](#); [Druckman, Peterson and Slothuus, 2013](#); [Guisinger and Saunders, 2017](#); [Haas and Khadka, 2020](#)).

Nearly all of these studies take the conventional approach of including as many confounders as possible as controls. For example, [Branton et al. \(2019\)](#) write: “In addition to the primary variables of interest, the model also includes several potentially confounding municipal-level social and economic demographic factors”, including the percentages of rural population, adults between the ages of 20 and 39, white voters, and female voters in each municipality, a measure of

²The exception is [Weintraub, Vargas and Flores \(2015\)](#), who find a non-monotonic relationship between exposure to violence in the Colombian conflict and support for President Santos in the 2014 election (which was centered around the negotiations with FARC). They find that municipalities with low violence *and* high violence were less likely to support Santos, while those with moderate levels of violence were more likely to support him.

³Going beyond observational work on the FARC vote as it took place, these results are corroborated by survey experiments that test the power of elite endorsements on attitudes toward peace ([Matanock and García-Sánchez, 2017](#); [Matanock and Garbiras-Díaz, 2018](#); [Matanock, Garbiras-Díaz and García-Sánchez, 2020](#)) as well as a 2014 survey in the field that found people's attitudes toward peace were shaped more by political preference than experience with violence ([Liendo and Braithwaite, 2018](#)).

government spending per municipality, and a measure of infant mortality per municipality (pg. 6). [Liendo and Braithwaite \(2018\)](#) note the importance of accounting for certain confounding traits because “the conflict has not affected civilians equally across the lines of ethnicity, gender, age, and socioeconomic conditions” (pg. 629). [Tellez \(2019b\)](#) includes “a number of controls in the base models that might confound inference”, which include respondent age, gender, monthly household income, education level, and level of trust in the national government, as well as municipal-level controls for support for the opposition party in 2010 (pgs. 1063–1065).

While adding these controls is a sensible starting point, it does not allow us to rule out confounding as the source of the observed relationships. Moreover, neither the statistical significance of these results nor their consistency across multiple studies tells us how sensitive they are in the face of potential confounders. Scholars often characterize claims from research of this type as “suggestive of” or “consistent with” claims that exposure to violence and political affiliation are causes of support for the peace deal. Regardless of such qualifications, these claims remain problematic: without further arguments, confounding may influence findings in either direction and with unknown magnitude, making them consistent not only with the claimed directional effects but with null effects or effects in the opposite direction.

Appendix B Sensitivity-based approaches

The use of formal sensitivity analyses in political science journals is, as noted, quite rare. Of the papers published in the *American Political Science Review*, *Journal of Politics*, and *American Journal of Political Science* in 2019, we coded 119 as relying upon observational research designs in which confounding would be a concern for their primary research question. This excludes principally field, survey, lab, and conjoint experiments. It includes “natural” and “quasi” experiments, as these are designs in which authors have reason to believe confounding is limited, but cannot rule it out entirely, making them ideal for sensitivity analyses. To be clear, a substantial fraction of papers (91, or 76%) engage in some form of robustness tests, typically procedures that attempt multiple specifications or sub-group analyses. While these have important uses, they do not quantify the amount of unobserved confounding that it would take to alter the conclusion. Thus, there is clearly a place for the sensitivity-based approach in observational social science research.

There have been many approaches to sensitivity since [Cornfield et al. \(1959\)](#), including [Rosenbaum and Rubin \(1983\)](#); [Heckman et al. \(1998\)](#); [Robins \(1999\)](#); [Frank \(2000\)](#); [Rosenbaum \(2002\)](#); [Imbens \(2003\)](#); [Brumback et al. \(2004\)](#); [Altonji, Elder and Taber \(2005a\)](#); [Hosman, Hansen and Holland \(2010\)](#); [Imai et al. \(2010\)](#); [Vanderweele and Arah \(2011\)](#); [Blackwell \(2013\)](#); [Frank et al. \(2013\)](#); [Dorie et al. \(2016\)](#); [Middleton et al. \(2016\)](#); [VanderWeele and Ding \(2017\)](#); [Oster \(2017\)](#); [Franks, D’Amour and Feller \(2019\)](#) and [Cinelli and Hazlett \(2020\)](#). From these approaches, we prefer the [Cinelli and Hazlett \(2020\)](#) one for several reasons.

First, we are interested in the sensitivity of estimates made using linear regression. Other important sensitivity methods are specialized to non-regression approaches, such as matching ([Rosenbaum, 2002](#)), which would be difficult here given that we examine two non-binary treatments. Sensitivity analyses for more general or non-parametric estimation procedures are possible, at the cost of requiring the user to make more elaborate characterizations of proposed confounding. For example, the “confounding function” approach ([Heckman et al., 1998](#); [Robins, 1999](#); [Brumback et al., 2004](#); [Blackwell, 2013](#)) generalizes across estimators, in cases with binary treatments. It requires the user to describe how the treated and untreated units would vary in their expectations of both the treated and untreated potential outcome, conditionally on the covariates. If investigators are willing to examine a linear model, as we and many others do, the

bias can instead be determined solely by the two parameters (or isomorphic variations thereof) deriving from omitted variable bias.

Further, among approaches to linear outcome models, [Imbens \(2003\)](#), [Carnegie, Harada and Hill \(2016\)](#) and [Dorie et al. \(2016\)](#) all require further assumptions beyond the two required ones, asking users to specify the distribution of the confounder as well as specifying the functional form of the treatment assignment mechanism. Relatedly, the approach taken by [Altonji, Elder and Taber \(2005a\)](#) and [Oster \(2017\)](#) employs a sensitivity parameter intended to reflect the relative predictive power of observables and unobservables in the selection (into treatment) process. However, this parameter is more complicated to interpret than it may at first seem, because it implicitly also requires contemplating how the observables and unobservables predict the outcome, as shown in [Cinelli and Hazlett \(2020\)](#).

Second, the availability of the $R_{Y \sim D|X}^2$ and particularly the RV as an interpretable, easy to convey, easy to compute sensitivity measure proves useful here.⁴ Additionally, this method for bounding/benchmarking corrects for issues in informal “benchmarking” practices employed in several other approaches. Informal benchmarking approaches such as those advocated in [Imbens \(2003\)](#); [Hosman, Hansen and Holland \(2010\)](#); [Dorie et al. \(2016\)](#); [Carnegie, Harada and Hill \(2016\)](#); [Middleton et al. \(2016\)](#); [Hong, Qin and Yang \(2018\)](#) aim to build intuition for the user by showing how a confounder “not unlike” an observed covariate in terms of its strength of relationship to the treatment and outcome would alter our conclusions. However, as shown in [Cinelli and Hazlett \(2020\)](#), those approaches can be misleading principally because even if confounding is assumed to be orthogonal to the included covariates, they become dependent when conditioning on the treatment. [Frank \(2000\)](#) largely avoids this concern by not conditioning on the treatment during benchmarking. Finally, the approach of [Altonji, Elder and Taber \(2005b\)](#), and subsequently [Oster \(2017\)](#), is technically correct, but as argued in [Cinelli and Hazlett \(2020\)](#), the parameterization of bias used there is problematic.

⁴A related approach is the E-value of [VanderWeele and Ding \(2017\)](#), which applies to relative risk estimates.

Appendix C Understanding sensitivity: Addressing common questions

C.1 Sensitivity analysis and causal inference frameworks

In general, one hopes to perform an analysis that rules out alternative explanations for why some D is related to some Y (other than that causal effect of D on Y). In the “selection on observables” set of approaches, the hope is that conditioning on some set of background covariates is sufficient to make the treatment independent of the potential outcomes, or equivalently block all backdoor paths from D to Y , or more colloquially, “render the $D - Y$ relationship unconfounded”. While using linear regression to perform the conditioning (rather than, say, matching or IPW weighting) imposes particular functional form restrictions, it is a common choice and a reasonable one in many circumstances.

The idea of sensitivity analysis in this realm is that, while achieving this unconfoundedness requires conditioning on $\{X, Z\}$, maybe you only conditioned on $\{X\}$, and you would like to know how missing the $\{Z\}$ might have harmed your answer. We emphasize that sensitivity tools can only speak to how the coefficient on D changes due to the inclusion of some hypothesized Z . Whether the investigator should be interested in the value of $\hat{\tau}$, i.e. the regression that includes Z , is up to the investigator and to the hypothesized variable Z . The conditions by which including Z identifies a causal quantity are well established, albeit with different terminology in different traditions. In the language of structural causal models (SCMs) or directed acyclic graphs (DAGs), we commonly require that $\{X, Z\}$ blocks all backdoor paths between D and Y , without opening new paths (by conditioning on colliders) or including post-treatment variables (Pearl, 2009). In the language of potential outcomes, we require that the potential outcomes at all levels of treatment are independent of the realized treatment assignment D_i conditionally on $\{X, Z\}$, i.e. $Y_i(d) \perp\!\!\!\perp D_i | \{X_i, Z_i\}$, often called (conditional) ignorability or selection on observables.

C.2 Sensitivity analysis for linear regression

In this paper we have engaged specifically with the sensitivity of linear regression results, as linear regression remains a common way of attempting to adjust for observables, not by fully conditioning on them in the non-parametric sense, but by adjusting for them through a model.

Again, this tool can only speak to how the coefficient (in a linear model) on D would change due to the inclusion of some hypothesized Z . In other words, even if one has the sensitivity

parameters in hand that describe $\{Z, X\}$ able to render D and Y unconfounded, one is still looking at a regression coefficient, and this may differ from quantities of interest such as the average treatment effect. Two chief concerns of this type are commonly raised. First, the regression coefficient offers only a “best linear summary” of the underlying relationship. Second, and consequently, because linear regression is “learning” more from areas of the data where there is variation in the treatment, the resulting estimate effectively up-weights observations from regions of X where there is more variability in the observed treatments (see e.g. Angrist and Pischke, 2008). If treatment effects are heterogeneous, and if the variation in effect size conditional on X is correlated with the variance in D conditional on X , then the resulting coefficient is a weighted average of the X -conditional ATEs, but not the usual average that weights according to $p(X)$. Such concerns may motivate investigators to employ approaches other than regression (such as weighting), and we would then suggest adopting sensitivity approaches suitable to those estimators.

Note, however, that the linearity presumption of linear regression is less problematic than may be expected when it comes to contemplating confounding strengths. In short, there is no requirement that Z 's relationship to either D or Y is “actually” linear. This is because whatever the nature of these relationships, the linear nature of the regression coefficient is only biased by the degree to which Z linearly explains D and Y , as indicated by the R^2 values. That is, the R^2 sensitivity parameters do not reflect true variance explainable if one had a non-parametric model, but simply the correlation squared, i.e. the linear component, because these are the only components that matter to the bias in the treatment coefficient. Further, Z can be conceptualized as any combination of omitted factors, and then the worst-case bias its omission causes is still characterized by these two R^2 terms. We direct readers to Cinelli and Hazlett (2020) for further details.

Another possibility of note is to conceptualize possible mis-specification in linear regressions as a form of omitted variable bias. For example, if one ought to have controlled for X_1^2 but only controlled for X_1 , the form of bias that may occur relates to how the part of X_1^2 not linearly explained by X_1 correlates with D and Y . One may therefore conceptualize this component as an omitted Z , allowing it to explain up to a certain amount of D and Y and determining the affect this has on the result. This is an active area of research.

C.3 Distinguishing partial R^2 from total R^2

To reinforce a point raised as a footnote in the text, it is useful to recall that partial R^2 (of a particular coefficient, i.e. the treatment) is distinct from the more familiar total R^2 of a model, or the “added R^2 ” from including an additional variable. Specifically, the partial R^2 of a treatment with the outcome (for example) is the fraction of residual variance in the outcome that it explains. If the covariates (other than treatment) can explain away 80% of the variance in the outcome, then the residual variance is 20%. If adding the treatment would then halve that, it would have a partial R^2 of 50%, not 10% (which would be the “added” R^2).

It is fundamentally the partial R^2 of the treatment with the outcome that determines the strength of confounding one would need to overturn it, as proven in [Cinelli and Hazlett \(2020\)](#). In models with low overall R^2 , the partial R^2 of the treatment with the outcome will be low in part because the treatment must not explain much variation in Y either. Even so, it is the partial R^2 of the treatment with outcome that drives the sensitivity, not the total R^2 . For models with high overall R^2 , we must be careful to remember that the high overall R^2 does not imply a high partial R^2 of the treatment. One could have, for example, fixed effects that push the overall R^2 up to something very high, say 95%. But the partial R^2 of the treatment is still free to vary widely as it reports only the amount additionally explained by the treatment. Moreover, even if there is very little residual variance leftover, confounding that does explain enough of that small variation can still have a strong impact on the bias depending on how it is related to the treatment.

C.4 Is the partial R^2 large or small?

Deciding whether a given partial R^2 value is “large” or “small” is the subject of additional analyses below given context-specific knowledge. However, to be clear on the meaning of these quantities, it is useful to recall that these partial R^2 values correspond literally to squared correlations. Thus, taking the square root of any R^2 allows interpretation on the usual correlation scale. That is, if a confounder is said to explain 0.32% of the residual variance in D conditional on X , for example, it means that $cor(D^{\perp X}, Z^{\perp X}) = \sqrt{0.0032} \approx 0.057$. In a context like ours where there is ample scope for unknown variables to influence treatment and outcome, this is a weak correlation indeed.

C.5 Is the RV large or small?

For the same reasons, we emphasize that the RV is only a report of how much confounding *it would take* to alter the conclusion. Since this is true regardless of *how much confounding there is*, there is no value of the RV that should be determined “good” or “bad”. It can only be used to indicate what must be argued from outside of the data to protect a conclusion. For example, an RV of 6% using observational data may make us very skeptical about the robustness of the result in the face of unobserved confounding, but an RV of 6% combined with a design that is meant to eliminate unobserved confounding (for example, a randomized control trial) might be high enough such that we do not think any confounding could explain that much residual variation.

C.6 Non-standard standard errors

While the “conventional” standard errors of a regression coefficient are constructed under assumptions that residuals are independent and homoskedastic, these assumptions may often be suspect. While a variety of more robust standard errors are available in common practice, integrating these into sensitivity analysis remains a challenge. With the tools of [Cinelli and Hazlett \(2020\)](#) employed here, this has no effect on the *point estimate* and its adjustment under confounding of a postulated strength. Hence the RV required to bring the coefficient to exactly zero ($q = 1$) is still correct. Contour plots show the bias-adjusted coefficients (such as Figure 1), as do the bounds shown on those plots that arise from benchmarking to observables. However, inferences that invoke the (adjusted) standard error will no longer be correct; as such, the $RV_{\alpha=0.05}$ value should not be used in these cases, nor should contour plots that show the bias-adjusted t-statistics rather than bias-adjusted coefficients.

C.7 Tutorials, examples, and software

A number of resources are available to demonstrate usage of this software in particular, and more importantly, to explore issues that arise in sensitivity analysis and to emphasize caveats and correct interpretation of the results.

- All analyses shown here were conducted using the `sensemkr` package in R, ([Cinelli, Ferwerda and Hazlett, 2020](#)), available at the Comprehensive R Archive Network (CRAN,

<https://cran.r-project.org/>). The documentation for this package (`?sensemakr`) provides basic example usage. Vignettes associated with the `sensemakr` package provide more extensive tutorials, including theoretical discussion, interpretational guidelines, and code examples. These can be found using `browseVignettes("sensemakr")` in R. The package is also available in Stata, by download from the Statistical Software Components (SSC) archive via `ssc install sensemakr, replace all`.

- The companion website, <https://carloscinelli.com/sensemakr/> provides news, links to presentations about the approach, and links to published and working papers.
- Replication materials for this paper can be found at <https://dataverse.harvard.edu/dataverse/jop>. Further, the `sensemakr` package for R also comes with data for this paper, which can be loaded using `data(colombia)`. The documentation for this dataset (`?colombia`) shows examples analyses replicating the analyses in this paper.
- Many of these analyses can be run within a web-browser (a “ShinyApp”), found at https://carloscinelli.shinyapps.io/robustness_value/.

Appendix D Models

D.1 Data

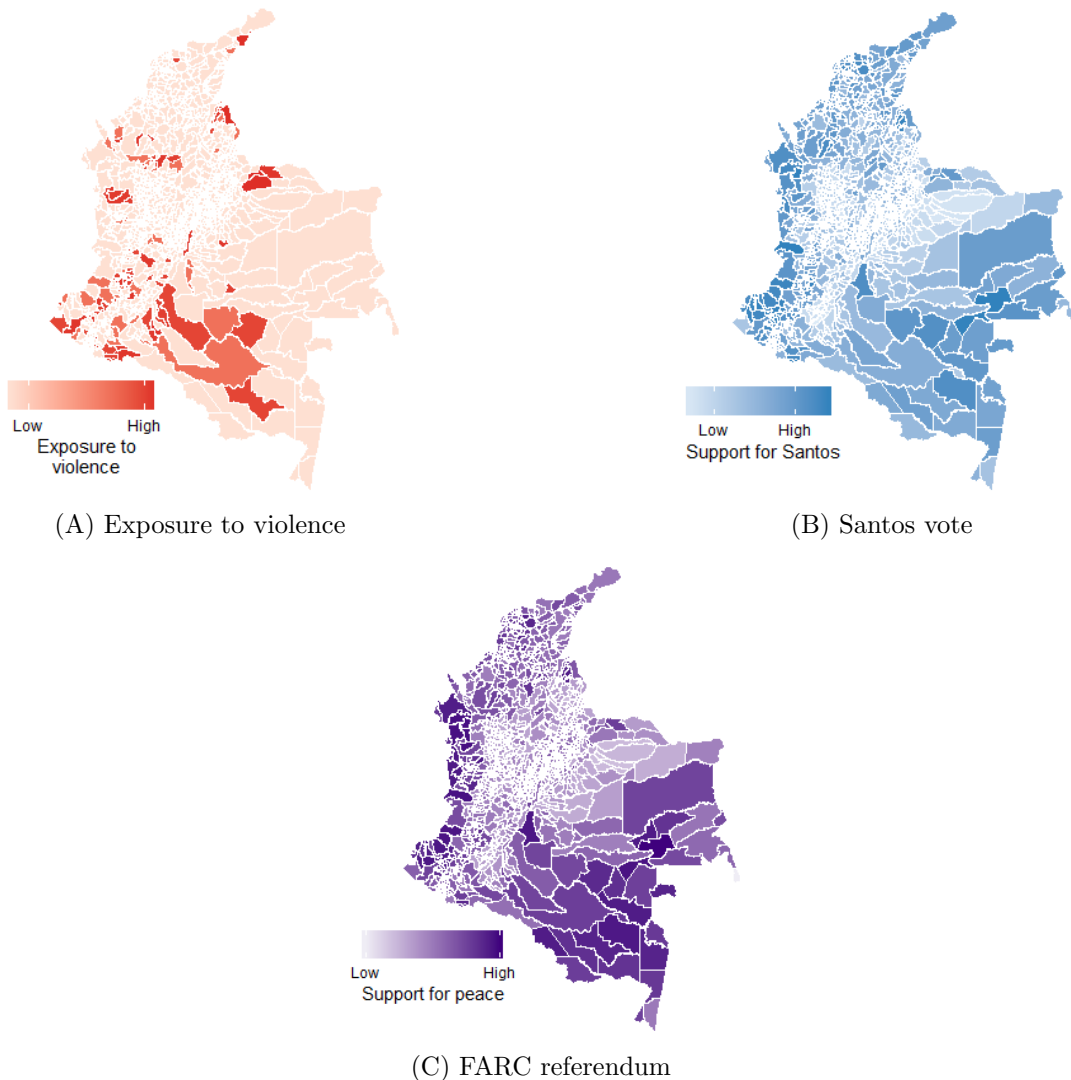
We measure FARC-related violence using incidences from the Global Terrorism Database, which has widespread coverage of events worldwide beginning in 1970. We chose this source because the database attributes attacks to particular groups, which is especially important in the Colombian context in which multiple guerilla groups were operating at the same time. We counted all FARC-related fatalities in each municipality per year, from any kind of attack, including, among others, murders, forced disappearance, and landmines. The exposure to violence measure is the cumulative number of FARC-related fatalities, grouped in years of five.

For the political affiliation hypothesis, we follow other studies and measure political affiliation as support for President Santos in the 2014 election, using results from the second round of the presidential elections. In some specifications, we use the 2010 vote share for President Santos (second round), since his position on negotiating with the FARC changed in 2012. All of our election results, including the results for the 2016 FARC referendum, are from the *Registraduría Nacional de Colombia*.

Figure A1 illustrates the spatial distribution of exposure to violence (Panel A), votes for President Santos in the 2014 election (Panel B), and votes in favor of the referendum for peace (Panel C) by municipality.⁵ Several municipalities in the southwestern part of the country (in the departments of Nariño and Cauca) were high on all three measures. In contrast, departments in the Andes Mountains are low on all three measures. Taken as a whole, exposure to violence (Panel A) seems to be somewhat similar in distribution to support for the FARC referendum (Panel C). The relationship between support for Santos (Panel B) and for the FARC referendum (Panel C) appears to be stronger still. We note that this strong visual relationship reflects that a large portion of the variance in the outcome is explained by support for Santos, which in the models appears as a $R_{Y \sim D|X}^2$ of nearly 60%. Recalling that this quantity is itself a useful sensitivity diagnostic, the strength of this relationship, even as seen graphically here, presages the robustness of the relationship between support for Santos and for the referendum.

⁵Maps were generated using the 'colmaps' package (Moreno, 2015) for creating maps of Colombia in the R statistical language (R Core Team, 2019).

Figure A1: Spatial Distribution of Key Variables



Note: Maps visualizing the municipal-level distribution of: (A) FARC-caused deaths, (B) vote share for Santos in the 2014 election, and (C) vote share in support of the FARC peace deal in the 2016 referendum.

D.2 Exposure to Violence

In addressing exposure to violence as an explanation for support, we consider two models. The first is a naive, direct comparison based on the simple model:

$$\text{Model 1: } Y_i = \beta_0 + \alpha(\text{Deaths}_{i,2011-2015}) + \epsilon_i \quad (4)$$

where Y_i is the proportion voting “Yes” in municipality i , and $\text{Deaths}_{i,2011-2015}$ is the number of deaths in municipality i committed by the FARC between 2011 and 2015. The coefficient α describes how the expected support for peace differs when we observe one additional death.

The second model takes the traditional approach of accounting for potentially worrying observed confounders by including them in the model as controls,

$$\text{Model 2: } Y_i = \beta_0 + \alpha(\text{Deaths}_{i,2011-2015}) + \beta_1(\text{Deaths}_{i,2006-2010}) + \beta_2(\text{Deaths}_{i,2001-2005}) + \beta_3(\text{Population}_i) + \beta_4(\text{GDPpc}_i) + \beta_5(\text{Santos } 2010_i) + \epsilon_i \quad (5)$$

where $\text{Deaths}_{i,2006-2010}$ and $\text{Deaths}_{i,2001-2005}$ are the number of deaths in municipality i in the corresponding time periods, Population_i is the total number of eligible voters, GDPpc_i is GDP per capita, and $\text{Santos } 2010_i$ is the vote share for President Santos in the 2010 election. We include the two lagged measures of violence, $\text{Deaths}_{i,2006-2010}$ and $\text{Deaths}_{i,2001-2005}$, as a means of accounting for areas that, for time-invariant reasons, routinely have higher or lower levels of violence. Indeed, we find that the “effect” of violence appears to fade over time, with only the most recent five years having a significant effect.⁶ Finally, we note that there are various additional ways of formulating these models – adding covariates or removing the lagged violence variables, for example – that can reduce the estimated effect of violence well below significance. The poor robustness of the model according to our analyses makes it unsurprising that we can so easily “ruin the result” by including different covariates. The sensitivity analysis would serve as a useful warning of the model’s fragility to alternative covariates, had we not been in a position to include and test them ourselves, as is the case for most readers of most papers. We proceed with the models favorable to the violence hypothesis for illustrative purposes. If even these models prove unable to withstand small degrees of confounding, alternative weaker models would generate even less persuasive results.

D.3 Political affiliation

We use the term “political affiliation” for this treatment but note that it is a shorthand: The key feature is not an individual’s support for a given leader, but whether the leader whom an individual supports publicly endorses the peace deal. In the present case, this simplifies to the

⁶In both models, for simplicity, we focus on violence in the 2011-2015 period as the treatment, because more recent violence more plausibly impacts attitudes. We note that GDP per capita is measured in 2013, which is post-treatment relative to violence occurring in 2011 and 2012; however, our assumption is that the effect of additional deaths at this level on GDP per capita is too small to be problematic.

question of whether a person supports President Santos. There are two ways to imagine the counterfactual outcome defining the treatment effect of interest. First, we can imagine how an individual might have voted in the referendum had they been loyal to a different leader but were otherwise unchanged. Though an individual's loyalties are associated with a variety of background factors, there is certainly enough room for variation that one can imagine this counterfactual. Alternatively, we can imagine an individual's vote in the FARC referendum, had their leader taken the opposite position. This is an easy possibility to entertain as well, since Santos was in fact against a peace deal with the FARC until 2012.

We estimate the coefficients in the model,

$$\text{Model 3: } Y_i = \beta_0 + \beta_1(\text{Santos } 2014_i) + \beta_2(\text{Deaths}_{i,2010-2013}) + \beta_3(\text{Elevation}_i) + \beta_4(\text{GDPpc}_i) + \beta_5(\text{Population}_i) + \epsilon, \quad (6)$$

where Y_i is the proportion voting "Yes" in municipality i , $\text{Santos } 2014_i$ is municipality vote share for Santos in the 2014 presidential election, and $\text{Deaths}_{i,2010-2013}$ is the total number of deaths due to FARC violence between 2010 and 2013 in municipality i .⁷ In this model we also control for total number of eligible voters (Population_i), as well as the mean elevation above sea level (Elevation_i) and GDP per capita of the municipality (GDPpc_i). These three control variables are chosen because they are unlikely to be affected by the treatment and because they are troubling potential confounders in as much as they could arguably be related to both political preferences and support for the FARC peace deal.

⁷Note that we use deaths between 2010–2013 as the pre-treatment covariate here, as the presidential election occurred in 2014.

Appendix E Sensitivity analyses from published regression results

Sensitivity analyses can aid readers and reviewers in assessing sensitivity even when authors may not have provided these analyses. The two summary statistics can be calculated using information in a regression table. First, $R_{Y \sim D | \mathbf{X}}^2$ requires only the t-statistic for the treatment coefficient and degrees of freedom from a regression, $R_{Y \sim D | \mathbf{X}}^2 = \frac{t_D^2}{t_D^2 + \text{dof}}$. Second, the *RV* can be calculated using published results: Let f_D be Cohen's partial f for the treatment variable, which can be obtained as $t_D / \sqrt{\text{dof}}$. Then $RV = 0.5(\sqrt{f_D^4 + (4f_D^2)} - f_D^2)$. Using these formulas, we can assess the sensitivity of published studies using linear regression.

For [Tellez \(2019b\)](#), we use the models reported in Appendix Table A5, which are the main regression models. We estimate the degrees of freedom to be approximately 4,200, as there are “roughly 4,200 observations” ([Tellez, 2019b](#), pg. 13). Across the three models shown there, the most favorable from a robustness point was Model (2) in Table A5 of that paper, which had an effect estimate of 0.22, and a standard error of 0.07, for a t-statistic of 3.14. This translates to a *RV* of 4.7%, indicating that a confounder explaining 4.7% of the residual variation in who is assigned to a conflict zone and in attitudes toward peace would be sufficient to eliminate the estimated effect. The effect would lose statistical significance at conventional levels with a confounder explaining just $RV_{\alpha=0.05}=1.8\%$ of these two residual variances. Finally, a confounder explaining all the remaining variation in the outcome need only explain 0.2% of who lives in a conflict zone in order to explain away the effect. Even a confounder explaining only 25% of the residual variation in the outcome would eliminate the estimated effect if it explains just 1% of residual variation in who is in a conflict zone.

For [Pechenkina and Gamboa \(2022\)](#), we use the results for the unmatched regression reported in Model 15 of the main results. The reported coefficient is 0.0423 and the standard error is 0.013, for a t-statistic of 3.25. We estimate about 1,008 degrees of freedom (1,052 reported observations, three main variables, eight controls, and 32 fixed effects). This results in an *RV* of 9.7%. The effect would lose statistical significance at conventional levels with a confounder explaining just $RV_{\alpha=0.05}=4\%$ of these two residual variances. Additionally, a confounder explaining all of the remaining variation in the outcome would only need to explain 1% of exposure to FARC violence in order to explain away the effect.

We took the values for [Krause \(2017\)](#) from from Model (2) of Table III, pg. 32. The

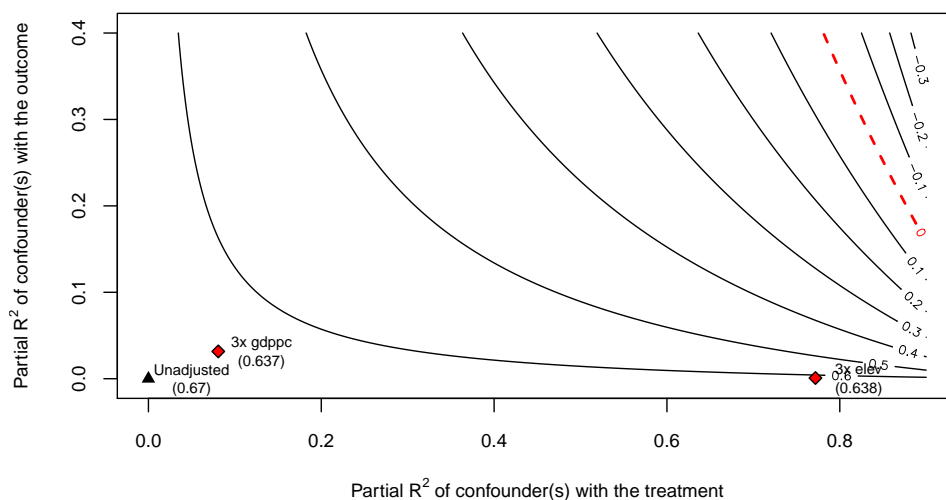
coefficient for Santos 2014 vote share in a similar model is 0.62, close to our estimate of 0.67.

The t-statistic of 45 together with 1,069 residual degrees of freedom would produce an *RV* of 72%, also similar to our own estimate of 68%. Note that this *RV* is an approximation because Krause (2017) reports Huber-White standard errors, whereas the formula for computing the *RV* from regression results calls for the conventional (spherical) standard error. If the conventional standard error were 20% larger than the Huber-White standard errors, for example, the *RV* would instead be 66%.

Appendix F Bounding for political affiliation

We know that it would take a strong confounder to overturn the results for political affiliation on voting. This does not mean, however, that such a confounder does not exist. The inability to “find” every confounder of interest is, of course, common in these studies. Yet, further progress can be made using the bounding approach described in the paper, transforming assumptions or statements about how confounding compares to observables into implied bounds on that confounding. One potential confounder is *GDP per capita*, which we expect might affect both treatment (political affiliation) and the outcome (support for the deal). Figure A2 shows the contour plot, to which we add a bound based on an assumption that “confounding is no more than three times ‘worse’ than GDP per capita,” in terms of the residual variation the confounder would need to explain in whom the voter supports and in support for peace (“3x gdppc”).⁸

Figure A2: Effect of unobserved confounding on estimate for political affiliation



Note: Contours showing adjusted regression coefficient on political affiliation, at levels of hypothesized confounders parameterized by the strength of relationship to the treatment (municipal vote for Santos in 2014) and the outcome (municipal vote for the FARC peace deal). The two bounds (“3 x gdppc” and “3x elev”) show the worst confounding that can exist, were we to assume that confounding is “no more than 3-times ‘worse’” than either GDP per capita or elevation. The dashed line indicates where the result would be eliminated.

The dashed line indicates the bound at which the result would be eliminated. A confounder three times as strong as that of *GDP per capita* would hardly reduce the estimate – from 0.67 to 0.64. Similarly, *Elevation* can also be used to formulate such a bounding assumption, as


⁸Note that we may wish to consider a k value for GDP per capita even higher than 3 to determine how far this robustness. As it turns out the maximum possible k value on this variable is 3.88. Such a proposed confounder would explain all of the residual variance of either the treatment or the outcome, and so a proposed confounder higher than this cannot exist. At $k = 3.88$, the point estimate still barely changes, to 0.63.

this relates to a wide variety of factors that may in turn relate to both political affiliation and preferences for peace with the FARC. Let us therefore assume that confounding is no more than three times “worse” than *Elevation* (“3x elev”). Again, the worst confounding that is possible under such an assumption would still hardly change the result, from 0.67 to 0.64. We emphasize that these bounds are linked to assumptions. In this case, while it is hard to imagine confounders more than three times “worse” than GDP per capita, we do not have sufficient knowledge (about what influences treatment or outcome) to ensure that no such confounder exists. We therefore do not regard these bounds as proof that our results are robust to confounding. Rather, they are “if-then” statements describing how strong confounding would have to be relative to these covariates in order to be problematic.

Next, we may be willing to make or probe assumptions — even pessimistic ones — about how much of the unexplained variance in the outcome could possibly be linked to confounding. We already know from the $R_{Y \sim D|X}^2$ value in Table 3 that a confounder explaining 100% of the residual variance of the outcome would need to explain 59% of the residual variance in political affiliation in order to overturn the result. Figure A3 provides an “extreme scenario” analysis that extends this reasoning. Each line shows what the adjusted effect estimate would be if confounding explains a proportion of the residual outcome (100%, 50%, or 30%) while explaining the proportion of treatment indicated by the horizontal axis. We see that less conservative confounding that explains 50% or 30% of the residual outcome variation would have to explain over 70% and 80% of the residual variation in political affiliation, respectively, to reduce the estimate to zero.

While we have determined that a powerful confounder would be required to alter our conclusion, it would be wrong to assume or imply that such a confounder does not exist without very convincing arguments as to why confounding is limited. To emphasize this point, we now discuss a potential confounder that one can imagine may be large enough to change our conclusions about political affiliation: attitudes toward the FARC. In particular, we could imagine that Colombians decided how to vote in the referendum based on how they felt about FARC *and* that attitudes toward FARC influenced their presidential vote in 2014, as the peace deal was a salient issue in the 2014 election. Santos announced his interest in negotiating with the FARC in 2012, and these negotiations became a major plank not just in Santos’ platform as he ran for reelection in 2014, but also in attacks made against him by Uribe and followers. The

Figure A3: Extreme scenario analysis



Note: Plot showing the extreme scenarios in which confounding explains 30%, 50%, and 100% of the residual outcome (municipal vote for the FARC peace deal) is explained by hypothesized confounding. The horizontal axis indicates a hypothesized proportion of residual variance explained in the treatment (municipal vote for Santos in 2014). A confounder explaining 100% of the residual variation in the outcome would need to explain 59% of the residual variance in the treatment (equivalent to $R_{Y \sim D|X}^2$), while a confounder explaining 50% of the residual variation in the outcome would need to explain about 75% of residual variation in the treatment. A confounder explaining 30% of the residual variation in outcome would need to explain over 80% of residual variation in the treatment to overturn the result.

2014 vote share for Santos was thus itself quite possibly affected by attitudes toward peace with the FARC. This constitutes a potentially powerful confounder that we imagine could explain more than 66% (to use $RV_{\alpha=0.05}$) of the residual variation in 2014 vote share and in support for the FARC referendum.

Ideally, we could control for such a confounder with some (pre-treatment) measure of FARC attitudes, but we have not been able to find such a variable.⁹ We note, however, that we can replace the 2014 vote share as our treatment with the 2010 vote share, in which peace with the FARC was not a particularly salient issue. Doing so, we still find that the relationship between political affiliation (as measured in 2010) and support for the 2016 peace deal would take large confounding to overturn ($R_{Y \sim D|X}^2 = 23\%$, $RV = 42\%$).

⁹While AmericasBarometer does ask respondents about their attitudes toward FARC and has for years, those data cover less than 6% of the municipalities in our data.

References

- Altonji, Joseph G., Todd E. Elder and Christopher R. Taber. 2005a. "An Evaluation of Instrumental Variable Strategies for Estimating the Effects of Catholic Schooling." *Journal of Human Resources* 40(4):791–821.
- Altonji, Joseph G., Todd E. Elder and Christopher R. Taber. 2005b. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113(1):151–184.
- Angrist, Joshua D and Jörn-Steffen Pischke. 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Bauer, Micheal, Christopher Blattman, Julie Chytilová, Joseph Henrich, Edward Miguel and Tamar Mitts. 2016. "Can War Foster Cooperation?" *Journal of Economic Perspectives* 30(3):249–74.
- BBC News. 2016. "Colombia referendum: Voters reject FARC peace deal." *BBC News* .
URL: <https://www.bbc.com/news/world-latin-america-37537252>
- Blackwell, Matthew. 2013. "A selection bias approach to sensitivity analysis for causal effects." *Political Analysis* 22(2):169–182.
- Branton, Regina, Jacqueline Demeritt, Amalia Pulido and James Meernik. 2019. "Violence, Voting & Peace: Explaining Public Support for the Peace Referendum in Colombia." *Electoral Studies* 61:1–13.
- Brumback, Babette A, Miguel A Hernán, Sebastien JPA Haneuse and James M Robins. 2004. "Sensitivity analyses for unmeasured confounding assuming a marginal structural model for repeated measures." *Statistics in medicine* 23(5):749–767.
- Carnegie, Nicole Bohme, Masataka Harada and Jennifer L Hill. 2016. "Assessing sensitivity to unmeasured confounding using a simulated potential confounder." *Journal of Research on Educational Effectiveness* 9(3):395–420.
- Cinelli, Carlos and Chad Hazlett. 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society, Series B* 82(1):39–67.

Cinelli, Carlos, Jeremy Ferwerda and Chad Hazlett. 2020. *sensemakr: Sensitivity Analysis Tools for Regression Models*. R package version 0.1.3.

URL: <https://CRAN.R-project.org/package=sensemakr>

Cornfield, Jerome, William Haenszel, E. Cuyler Hammond, Abraham M. Lilienfeld, Michael B. Shimkin and Ernst L. Wynder. 1959. "Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions." *Journal of National Cancer Institute* (23):173–203.

Dávalos, Eleonora, Leonardo Fabio Morales, Jennifer S. Holmes and Liliana M. Dávalos. 2018. "Opposition Support and the Experience of Violence Explain Colombian Peace Referendum Results." *Journal of Politics in Latin America* 10(2):99–122.

Dorie, Vincent, Masataka Harada, Nicole Bohme Carnegie and Jennifer Hill. 2016. "A Flexible, Interpretable Framework for Assessing Sensitivity to Unmeasured Confounding." *Statistics in Medicine* 35(20):3453–3470.

Druckman, James, Erik Peterson and Rune Slothuus. 2013. "How Elite Partisan Polarization Affects Public Opinion Formation." *American Political Science Review* 107(1):57–79.

Fabbe, Kristin, Chad Hazlett and Tolga Sınmazdemir. 2019. "A persuasive peace: Syrian refugees' attitudes towards compromise and civil war termination." *Journal of Peace Research* 56(1):103–117.

Frank, Kenneth A. 2000. "Impact of a confounding variable on a regression coefficient." *Sociological Methods & Research* 29(2):147–194.

Frank, Kenneth A, Spiro J Maroulis, Minh Q Duong and Benjamin M Kelcey. 2013. "What would it take to change an inference? Using Rubin's causal model to interpret the robustness of causal inferences." *Educational Evaluation and Policy Analysis* 35(4):437–460.

Franks, Alex, Alex D'Amour and Avi Feller. 2019. "Flexible Sensitivity Analysis for Observational Studies without Observable Implications." *Journal of the American Statistical Association* 115(1532):1730–1746.

Gallego, Jorge, Juan D Martínez, Kevin Munger and Mateo Vásquez-Cortés. 2019. "Tweeting for peace: Experimental evidence from the 2016 Colombian Plebiscite." *Electoral Studies* 62:102072.

- Guisinger, Alexandra and Elizabeth Saunders. 2017. "Mapping the Boundaries of Elite Cues: How Elites Shape Mass Opinion across International Issues." *International Studies Quarterly* 61(2):425–441.
- Haas, Nicholas and Prabin B. Khadka. 2020. "If They Endorse It, I Can't Trust It: How Outgroup Leader Endorsements Undercut Public Support for Civil War Peace Settlements." *American Journal of Political Science* 64(4):982–1000.
- Hazlett, Chad. 2020. "Angry or Weary? How Violence Impacts Attitudes Toward Peace among Darfurian Refugees." *Journal of Conflict Resolution* 64(5):844–870.
- Heckman, James, Hidehiko Ichimura, Jeffrey Smith and Petra Todd. 1998. "Characterizing Selection Bias Using Experimental Data." *Econometrica* 66(5):1017–1098.
- Hong, Guanglei, Xu Qin and Fan Yang. 2018. "Weighting-Based Sensitivity Analysis in Causal Mediation Studies." *Journal of Educational and Behavioral Statistics* 43(1):32–56.
- Hosman, Carrie A, Ben B Hansen and Paul W Holland. 2010. "The Sensitivity of Linear Regression Coefficients' Confidence Limits to the Omission of a Confounder." *The Annals of Applied Statistics* 4(2):849–870.
- Hurtado-Parrado, Camilo, Myriam Sierra-Puentes, Mohammed El Hazzouri, Alexandra Morales, Diana Gutiérrez-Villamarín, Laura Velásquez, Andrea Correa-Chica, Juan Carlos Rincón, Karen Henao, Juan Gabriel Castañeda et al. 2019. "Emotion regulation and attitudes toward conflict in Colombia: Effects of reappraisal training on negative emotions and support for conciliatory and aggressive statements." *Frontiers in Psychology* 10:1–9.
- Imai, Kosuke, Luke Keele, Teppei Yamamoto et al. 2010. "Identification, inference and sensitivity analysis for causal mediation effects." *Statistical science* 25(1):51–71.
- Imbens, Guido W. 2003. "Sensitivity to exogeneity assumptions in program evaluation." *The American Economic Review* 93(2):126–132.
- Krause, Dino. 2017. "Who Wants Peace? The Role of Exposure to Violence in Explaining Public Support for Negotiated Agreements: A Quantitative Analysis of the Colombian Peace Agreement Referendum in 2016." Unpublished Master's Thesis, Uppsala University.

- Levendusky, Matthew. 2010. "Clearer Cues, More Consistent Voters: A Benefit of Elite Polarization." *Political Behavior* 32(1):111–131.
- Liendo, Nicolás and Jessica Maves Braithwaite. 2018. "Determinants of Colombian Attitudes toward the Peace Process." *Conflict Management and Peace Science* 35(6):622–636.
- Lupia, Arthur. 1994. "Shortcuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections." *American Political Science Review* 88(1):63–76.
- Masullo, Juan and Davide Morisi. 2019. "Risks and opportunities of direct democracy: the effect of information in Colombia's peace referendum." *Politics and Governance* 7(2):242–267.
- Matanock, Aila and Miguel Garcia-Sánchez. 2017. "The Colombian Paradox: Peace Processes, Elite Divisions & Popular Plebiscites." *Daedalus* 146(4):152–166.
- Matanock, Aila and Natalia Garbiras-Díaz. 2018. "Considering Concessions: A Survey Experiment on the Colombian Peace Process." *Conflict Management & Peace Science* 35(6):637–655.
- Matanock, Aila, Natalia Garbiras-Díaz and Miguel García-Sánchez. 2020. "Elite Cues and Attitude Formation in Post-Conflict Contexts." Empirical Studies of Conflict Project Working Paper No. 19.
URL: <https://esoc.princeton.edu/WP19>
- Meernik, James. 2019. "Violence and Reconciliation in Colombia: The Personal and the Contextual." *Journal of Politics in Latin America* 11(3):323–347.
- Middleton, Joel A, Marc A Scott, Ronli Diakow and Jennifer L Hill. 2016. "Bias amplification and bias unmasking." *Political Analysis* 24(3):307–323.
- Moreno, Daniel. 2015. *colmaps: Colombian maps*. R package version 0.0.0.9515.
- Nicholson, Stephen. 2012. "Polarizing Cues." *American Journal of Political Science* 56(1):52–66.
- Oster, Emily. 2017. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics* 37(2):187–204.
- Pearl, Judea. 2009. *Causality*. Cambridge University Press.

- Pechenkina, Anna O. and Laura Gamboa. 2022. "Who Undermines the Peace at the Ballot Box? The Case of Colombia." *Terrorism & Political Violence* 34(2):197–217.
- Petersen, Roger and Sarah Zuckerman Daly. 2010. Anger, Violence, and Political Science. In *International Handbook of Anger: Constituent and Concomitant Biological, Psychological, and Social Processes*, ed. Charles Spielberger Michael Potegal, Gerhard Stemmler. Springer.
- R Core Team. 2019. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
URL: <https://www.R-project.org/>
- Revelo, Diana Rico and Cecilia Emma Sotttilotta. 2020. "Barriers to Peace? Colombian Citizens' Beliefs and Attitudes Vis-à-Vis the Government-FARC-EP Agreement." *Studies in Conflict & Terrorism* pp. 1–22.
- Robins, James M. 1999. "Association, causation, and marginal structural models." *Synthese* 121(1):151–179.
- Rosenbaum, Paul R. 2002. Observational Studies. In *Observational Studies*. Springer pp. 1–17.
- Rosenbaum, Paul R and Donald B Rubin. 1983. "Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome." *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 212–218.
- Tellez, Juan Fernando. 2019a. "Peace Agreement Design and Public Support for Peace: Evidence from Colombia." *Journal of Peace Research* 56(6):827–844.
- Tellez, Juan Fernando. 2019b. "Worlds Apart: Conflict Exposure and Preferences for Peace." *Journal of Conflict Resolution* 63(4):1053–1076.
- Vanderweele, Tyler J. and Onyebuchi A. Arah. 2011. "Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders." *Epidemiology (Cambridge, Mass.)* 22(1):42–52.
URL: <http://dx.doi.org/10.1097/ede.0b013e3181f74493>
- VanderWeele, Tyler J and Peng Ding. 2017. "Sensitivity analysis in observational research: introducing the E-value." *Annals of internal medicine* 167(4):268–274.

Weintraub, Michael, Juan Vargas and Thomas Flores. 2015. "Vote Choice and Legacies of Violence: Evidence from the 2014 Colombian Presidential Elections." *Research and Politics* 2(2):2053168015573348.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge University Press.